Online Appendix of Intertemporal Substitution and New Car Purchases *

Adam Copeland[†]

December 12, 2013

Abstract

This article presents a dynamic demand model for motor vehicles. This approach accounts for the change in the mix of consumers over the model year and measures consumers' substitution patterns across products and time. I find intertemporal substitution is significant; consumers are more likely to change the timing of their purchase in reaction to a price increase rather than buy another vehicle in the same period. Further, I find automakers' use of large cash-back rebates at the end of the model year, although boosting overall sales, induces large numbers of consumers to delay their purchases and so pay lower prices.

Key words: discrete-choice demand estimation, automobiles, dynamics
JEL classification: D12, C61, L62

*I thank Victor Aguirregabiria, Ana Aizcorbe, Richard Crump, Liran Einav, Ron Goettler, Gautam Gowrisankaran, Ryan Greenaway-McGrevy, Igal Hendel (editor), Kathleen Johnson, Ahmed Khwaja, Donghoon Lee and Eugenio Miravete as well as two anonymous referees for their comments and suggestions. I started writing this article while at the U.S. Bureau of Economic Analysis. The views expressed here are those of the author and do not necessarily reflect the position of the Bureau of Economic Analysis, the Department of Commerce, the Federal Reserve Bank of New York, or the Federal Reserve System.
[†]Federal Reserve Bank of New York; e-mail: adam.copeland@ny.frb.org

1 On-line Appendix

In this appendix, I describe the computational details used to produce the results found in "Intertemporal Substitution and New Car Purchases". This estimation approach is a variation of that described in Gowrisankaran and Rysman (2012).

Some notation (see the article for more details):

- $\theta \equiv \{\gamma, \zeta, \alpha, \sigma\}$: parameters to be estimated,
- **X**: observed non-price vehicle characteristics,
- ξ: unobserved vehicle characteristic,
- **p**: vehicle prices,
- **Z**: instruments for price,
- **C**: monthly dummy variables,
- **y**: household income,
- ν: households' taste for characteristics
- **I**: a 4 element vector of the observed average income of new vehicle purchasers by quarter (Aizcorbe, Bridgman, and Nalewaik (2010))

The consumer's new vehicle purchasing problem

As detailed in equation 3 of the article, consumer i's indirect utility from purchasing model j in month t of the model year is given by:

$$u_{ijt} = \mathbf{X}_{jt}\gamma + \sum_{s=1}^{11} 1_{C_t=s}\zeta_s + \xi_{jt} - \frac{\alpha}{y_i}p_{jt} + \sum_{k=1}^{K} \sigma_k \nu_{ik} x_{jkt} + \varepsilon_{ijt},$$

where ε is an iid error term distributed type 1 extreme value and $x_{jkt} \in \mathbf{X}_{jt}$ is the kth observable characteristic of product j at time t. Collecting terms, we can rewrite this as

$$u_{ijt} = \lambda_{jt} + \mu_{ijt} + \varepsilon_{ijt}.$$

where

$$\lambda_{jt} = \mathbf{X}_{jt}\gamma + \sum_{s=1}^{11} 1_{C_t=s}\zeta_s + \xi_{jt} \quad \text{and,}$$

$$\mu_{ijt} = \frac{-\alpha}{y_i} p_{jt} + \sum_{k=1}^{K} \sigma_k \nu_{ik} x_{jkt}.$$

Let ε_{i0t} denote consumer i's indirect utility from the outside option at time t and assume it is iid and distributed type 1 extreme value. We can write the consumer i's value function in the last month of the model year as

$$V_{iT} = \max\left\{\varepsilon_{i0T}, \max_{j\in J}\left\{u_{ijT}\right\}\right\},$$

where J is the set of vehicles available for purchase. For the months preceding the last month of the model year, $t = 1, 2, \ldots T-1$, the consumer's value function is

$$V_{it} = \max\left\{\varepsilon_{i0t} + \beta E[V_{i,t+1}], \max_{j\in J}\left\{u_{ijt}\right\}\right\},$$

where β is the discount factor (which I set to $0.95^{1/12}$). E denotes the expectation taken over next period's choice set, vehicle characteristics, prices, monthly dummies, and ε. I assume consumers have perfect foresight over the evolution of the choice set, vehicle characteristics, prices and the monthly dummies. Given the distributional assumptions on both ε_{ijt} and ε_{i0t}, we know that

$$E[V_{iT}] = \log\left(\sum_{j\in J} \exp(\lambda_{jT} + \mu_{ijT}) + \exp(0)\right). \tag{1}$$

For $t = 1, 2, \ldots, T-1$, we have

$$E[V_{it}] = \log\left(\sum_{j\in J} \exp(\lambda_{jt} + \mu_{ijt}) + \exp(\beta E[V_{i,t+1}])\right). \tag{2}$$

We use this framework to predict when households make purchases. Let d_{it} denote consumer i's purchase decision in period t, where $d_{it} = 0$ means no purchase. Because of the distributional assumptions on ε_{ijt} and ε_{i0t}, consumer i's probability of purchasing a vehicle j at time t is equal to

$$Pr(d_{it} = j) = \frac{\exp(\lambda_{jt} + \mu_{ijt})}{\sum_{g \in J} \exp(\lambda_{gt} + \mu_{igt}) + \exp(\beta E[V_{i,t+1}])}, \tag{3}$$

By aggregating over consumers, we can compute a predicted market share. In this framework, aggregating over consumers is complicated by the fact that the mass of potential consumers available over the model year evolves. This is because I assume that households leave the market after purchasing a new vehicle. To account for this evolution, I assume that a consumer represents a type of household. In the first month of the model year, each consumer type has measure 1. If in month 1 consumer i has a probability of purchasing any vehicle (any inside good) of x percent, then the measure of consumers of type i in month 2 is equal to (1-x). Formally, letting m_{it} denote the measure of consumer i, then

$$m_{i1} = 1, \tag{4}$$
$$m_{it} = m_{i,t-1} * \big(1 - \sum_{j \in J} Pr(d_{i,t-1} = j)\big) \text{ for } t = 2, 3, \ldots, T. \tag{5}$$

Given m_{it}, the predicted market share of a vehicle j in month t is

$$s_{jt} = \frac{\sum_{i=1}^{N} Pr(d_{it} = j)\, m_{it}}{\sum_{i=1}^{N} m_{it}}. \tag{6}$$

In practice, I used 200 consumer types to estimate the model (i.e. $N = 200$).

The moments

There are two sets of moments. The first are the usual moments constructed from the unobserved vehicle characteristic. They are

$$G_1(\theta) = \mathbf{Z}'\xi(\theta)), \tag{7}$$

where $\mathbf{Z}$ are the exogenous variables used to instrument for price. The second set of moments are constructed from the average income of new vehicle purchasers over the model year. Letting $I(\theta)$ denote the model's prediction of the mean income of new vehicle purchases for each quarter of the model year, the second set of moments are

$$G_2(\theta) = I(\theta) - \mathbf{I}. \tag{8}$$

The GMM loss criterion is

$$\begin{bmatrix} G_1(\theta) \\ G_2(\theta) \end{bmatrix}' \mathbf{W}^{-1} \begin{bmatrix} G_1(\theta) \\ G_2(\theta) \end{bmatrix} \tag{9}$$

where $\mathbf{W}$ is a block diagonal matrix. The block of elements of $\mathbf{W}$ associated with G_1 are equal to $\mathbf{Z}'\xi(\hat{\theta})\xi(\hat{\theta})'\mathbf{Z}$, where $\hat{\theta}$ is a consistent estimate of the true θ. The block of elements of $\mathbf{W}$ associated with G_2 are equal to a diagonal matrix, where the elements of the diagonal are equal to the standard error of the mean income of new vehicle purchasers by quarter. I used the standard errors of the means to provide relative weights across the different income means. This approach to constructing the weighting matrix results in consistent, but not efficient, parameter estimates.[1]

As a robustness check, I also estimated the model with an alternative weighting matrix. The portion of the weighting matrix associated with the first set of moments remains equal to $\mathbf{Z}'\xi(\hat{\theta})\xi(\hat{\theta})'\mathbf{Z}$. Let $D(\hat{\theta})$ be the difference between the model's predictions of average income per quarter and the data. The portion of the weighting matrix associated with the second set of moments is equal to the standard errors of the means times the square of $D(\hat{\theta})$. The estimated parameters using this alternative weighting matrix are only trivially different from those reported in the article.

The estimation algorithm

In this subsection, I describe how to estimate the parameters of the model. The goal of the optimization routine is to minimize the moments described in the previous section. To help

[1] I follow the usual approach to obtain a consistent estimate of θ. I estimate the model, minimizing the GMM loss criterion function with the following weighting matrix. The portion of weighting matrix associated with the first set of moments is set equal to $\mathbf{Z}'\mathbf{Z}$. The portion of the weighting matrix associated with the second set of moments is equal to a diagonal matrix, wherein the diagonal elements are equal to the standard errors of the means. The resulting parameter estimates, $\hat{\theta}$, are consistent estimates of the true value of θ.

follow this description, please refer to the accompany schematic (see figure 1).

1. The algorithm starts with an initial guess of a subset of the parameters, (α, σ), random draws of household's income, y, and of household's tastes for vehicle characteristics, ν. An initial guess of $E[\mathbf{V}]$ (consumers' option value of waiting, as defined in equations 1 and 2) is also required (box 1 in figure 1).

2. The next portion of the algorithm is a loop whose goal is to arrive at estimates of ξ_{jt} (the unobserved vehicle characteristic) and $E[\mathbf{V}]$ which are consistent with one another (boxes 2,3,4 in figure 1).

3. To this end, the next step is to use the contraction mapping described in Berry, Levinsohn, and Pakes (1995) to recover ξ_{jt} (box 2 in figure 1). The ξ_{jt} are computed so that the model's prediction of vehicle j's sales at time t match the data. Note that the model's prediction of sales requires knowing consumers' $E[\mathbf{V}]$, which, at this point, we have only guessed. For this model, the ξ's must be computed sequentially, starting from the first month of the model year. This is because the set of potential consumers in month t depend upon purchase decisions in month $t-1$ (see equation 4). Hence, first I recover ξ_{j1} for all j knowing that $m_{i1} = 1$ for all i. I then compute m_{i2} for all i and recover ξ_{j2} for all j, and so on.

4. I now turn to updating $E[\mathbf{V}]$ (box 3 in figure 1). As detailed in equations 1 and 2, consumers' value functions, in expectation, are a function of ξ_{jt} and μ_{ijt}. In the previous step, I recovered ξ_{jt}, and μ_{ijt} is a function of the parameters (α, σ) and the random draws ν. Via backwards induction I can quickly calculate updated values of $E[V_{it}]$ for all (i,t).

5. I determine whether ξ_{jt} and $E[\mathbf{V}]$ are consistent by seeing by how much the updated values of $E[\mathbf{V}]$ have changed from its previous values (box 4 in figure 1). Formally, the mean absolute difference between these two vectors needs to be less than 1e-12.

6. If this condition is not met, I compute new values of ξ_{jt} using the updated values of $E[\mathbf{V}]$ and repeat the steps above. Although it has not been proven that this process should converge (i.e. that subsequent updated values of $E[\mathbf{V}]$ will change by smaller and smaller amounts), in practice I only experienced problems along this dimension for parameter values far from the estimated values.

7. If the convergence condition is met, I move to the next portion of the algorithm—evaluating the GMM loss criterion (box 5 in figure 1). Note that given the model's predictions about consumer *i*'s purchasing decisions, it is straightforward to calculate the model's predictions of the mean income of new vehicle purchases.

8. If the GMM loss criterion is minimized, then the estimation algorithm stops and reports results. If the GMM loss criterion is not minimized, then new values of (α, σ) are chosen (box 6 in figure 1) and the process is repeated.

I used a non-derivative-based simplex algorithm to find the parameter vector which minimizes the GMM criterion function. An alternative approach is laid out in Dube, Fox, and Su (2012). They report that their algorithm is faster and avoids some numerical issues associated with the looping structure described above.

Figure 1: Schematic of the estimation algorithm

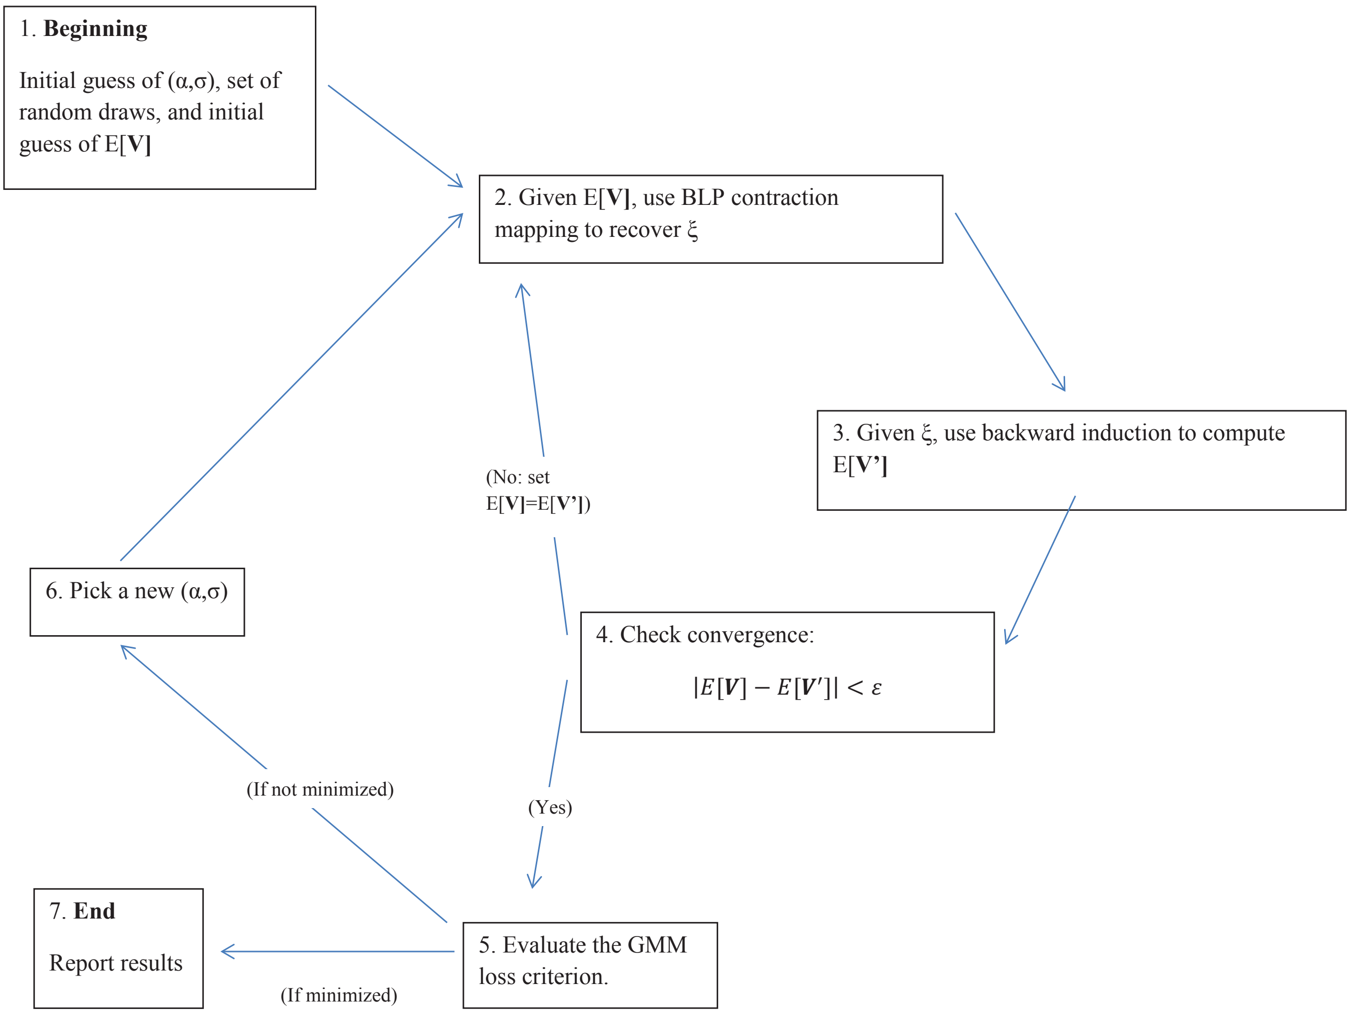

Note: This schematic was inspired by figure 3 in Schiraldi (2011).

References

AIZCORBE, A., B. BRIDGMAN, AND J. NALEWAIK (2010): “Heterogeneous Car Buyers: A stylized fact,” *Economic Letters*, 109, 50–53.

BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): “Automobile Prices in Market Equilibrium,” *Econometrica*, 63, 841–890.

DUBE, J.-P., J. FOX, AND C.-L. SU (2012): “Dynamic Discrete Choice Random Coefficients Demand Estimation,” *Econometrica*, 80(5), 2231–2267.

GOWRISANKARAN, G., AND M. RYSMAN (2012): “Dynamics of Consumer Demand for New Durable Goods,” *Journal of Political Economy*, 120(6), 1173–1219.

SCHIRALDI, P. (2011): “Automobile Replacement: a Dynamic Structural Approach,” *Rand Journal of Economics*, 42(2), 266–291.